

CAS Operations Integration Planning

Bill Boroski

CAS Review
March 1-2, 2004

Outline

- Objectives
- Process Elements
 - Handoff to operations
 - Integration into DP factory
 - Database hosting and hardware systems support
 - Improved Collaboration access to data
 - Public release schedules
- Resource Availability and Requirements
- Summary

Objectives

- Objectives
 - Complete the hand off of the CAS loading operation from JHU to the Fermilab DP group.
 - Integrate the CAS loading operation into the Fermilab DP factory so that CAS loading is performed a production manner (reliable and repeatable);
 - Serve “finished” CAS databases and the SkyServer to the SDSS collaboration and the general public from Fermilab in a robust and reliable manner.
 - Reduce the length of time it takes to get processed data loaded into the CAS and into the hands of the collaboration.
 - Release the DR3, DR4 and Final Release CAS to the general astronomy community according the NSF-approved data release schedule.

Handoff to Operations

- Requires a stable data model
 - Data model changes have occurred for each data release to date.
 - Propagating data model changes through all phases of CAS loading has not occurred smoothly.
 - Every data release to date has been done in JIT mode and under extreme schedule pressure.
 - Under these conditions, transfer of knowledge from developers to operations has been marginally successful at best.
 - Data model changes are planned for the next two releases (DR3 and DR4).
 - Initial DR3 load will use DR2 data model; provides an opportunity for the DP group to get a grasp on the loading operation.
 - DR3 CSV creation has gone smoothly. Solid evidence that stability will help the transition.
 - Final DR3 load must be finished by July 1. Time between now and mid-May can be used to implement and test final DR3 spectro data model changes at JHU and propagate these into the DP factory

Handoff to Operations (2)

- Stable data model is unlikely any time soon.
 - As a project we need to manage this better.
 - We will not re-load DR3 in the factory until we can successfully and repeatedly load DR3 data using the DR2 data model. This baselines the loading operation.
 - DR3 reload can be tested and debugged at JHU, then transferred to Fermilab.
 - DR4 re-processing and CAS loading schedule will be developed that factors in adequate testing time at JHU and in the DP factory.

Handoff to Operations (3)

- Requires adequate documentation
 - Ani has worked very hard to document loading steps and to keep documentation up-to-date.
 - A significant amount of documentation exists
 - The committee is encouraged to review what we have.
 - Critical to limit data model changes.
 - Provides more time to shore up documentation. Hard to nail down a moving target.
- Operations group can help develop documentation as they gain experience with a stable system
 - Has been happening at a limited level by John I. and Vijay.
 - Engaging the receiving group in documentation development follows model for developing/improving observing systems documentation at APO.
 - Not ideal, but given limited resources, is necessary.

Handoff to Operations (4)

- Transfer of Knowledge
 - Impeded by geographical dispersion, limited Windows and SQL Server knowledge within the DP factory, changing data models, schedule pressures.
 - To address these challenges:
 - We won't move forward in factory with DR3 reload until we have the DR3 initial load process understood.
 - We are actively bringing in more resources to support operations at Fermilab.
 - The DR3 public release schedule gives us some breathing room to get the CAS loading operation running more smoothly in the factory.
 - We will have more face-to-face meetings between the development and operations groups.

Handoff Plan

- No new CSV creation or sqlLoader code installed at Fermilab until we have the loading operation understood using the DR2 data model and associated CAS code.
- Once DR2 is released to public and DR3 pre-load is released to the collaboration, updating the loading documentation will be given highest priority.
 - Documentation will be improved until procedures accurately describe the loading process.
 - Documentation status will be reviewed at weekly meetings.
- Bring more experienced resources into the loading operation at Fermilab to increase level of expertise in Windows and SQL Server.
 - Support commitment will be formalized with CD/CSS/CSI department
 - Proposal will be developed to increase level of effort in DP factory.
- More frequent face-to-face meetings will be scheduled to improve communication and focus more attention on knowledge transfer.

Integration into the Factory

- CSV File Creation
 - Most integrated of the CAS processes
 - Creation of Best and Target CSVs for DR3 went relatively smoothly
 - Error-checking implemented to improve reliability and robustness
 - The Committee is requested to review what has been put in place and comment on applicability to other areas.
- New Hardware Setup Procedure
 - Documentation exists but is not complete; still being refined.
 - Setting up two new 6TB SATA servers for DR3 took ~ 2 weeks.
 - Additional assistance from CD/CSS/CSI department will benefit this effort.
 - Documentation and expert help will reduce time required to bring new machines on-line. Especially true given long periods (4-6 mos.) between new machine purchases.

Integration into the Factory (2)

- CAS Loading Operations
 - To reiterate:
 - We will work with DR2 data model until loading process is understood by operations group and is repeatable. DR3 pre-load is the pilot project.
 - sqlLoader code will remain stable (no new features, only bug fixes).
 - Documentation will be continuously improved as we move forward. Joint effort between development and operations groups.
 - We will increase face time between development and operations groups.
- SkyServer Front-end and Web Servers
 - This area has not yet been explored in any detail.
 - Implementation and connections have been performed by the development group to date (especially true for SkyServer front-end).
 - Transfer to operations will be on back-burner until loading process is understood; until then, responsibility remains with development group.
 - Documentation exists but not fully evaluated and used by the operations group.

Factory Integration Planning

- Use DR3 pre-load to improve procedures, increase internalize knowledge of the loading process, and increase confidence.
 - Do not introduce new code versions until loading with the DR2 data model is repeatable and reliable.
- Implement QA checks in loading process using experience gained in implementing QA checks in CSV creation process.
- Work closely with developers to learn how to install, configure and manage SkyServer front-end.
 - Face-to-face meetings will be used to more effectively transfer knowledge.

CAS /SkyServer Hosting Issues

- Primary copies of public EDR and DR1 SkyServers are currently being served from JHU. Backup copies are hosted at Fermilab.
 - Public DR1 SkyServer is averaging 1.5 million hits
 - 1.3 million SQL queries executed in February 2004.
 - DR1 SkyServer usage to date = 36.6 million hits; 2.7 million SQL queries
- Goal is to transition primary public site to Fermilab
 - We are considering this for public DR2 release in mid-March.
 - Requires implementation of automatic server switch-over, systems monitoring, and on-call support.
 - Fall-back is to host primary site at JHU for DR2; move primary site to Fermilab for DR3.
- Need to transfer knowledge from development team to operations team.
- Anticipated technical support and assistance from the CD/CSS/CSI group is a key element in providing solid performance and reliability.
 - Experience and expertise running high availability systems.
 - Experience and expertise running Windows-based systems.

CAS / SkyServer Hosting Plan

- We do not have a concrete plan for transferring primary hosting responsibilities at this time.
- We need to arrange meetings between the development and operations team to understand full scope of work involved in hosting primary site at Fermilab
 - Need to consider education and outreach portions of the SkyServer
 - Need to clearly define roles, responsibilities, and access rights.
- Develop a transfer plan and timeline based on discussions between JHU development team and Fermilab operations team (EAG and CD/CSS/CSI).

Better Data Access to SDSS Collaboration

- Collaboration currently gets data loaded into CAS ~3 months prior to public release.
- Original plan was to provide collaboration access 12-18 months in advance.
 - Technically, collaboration has access within one month of data acquisition in the form of flat data files.
 - Not good enough.
- Releasing DR3 pre-load to collaboration in mid-March provides 6 months of pre-release access.
- Releasing DR4 to collab by December 2004 provides 6 months of pre-release access.
 - We will do what we can to move this up, within reason.
- We need to develop a realistic work plan and schedule for DR3-and-beyond data releases.

Data Release Schedules

- Data Release 1 occurred 6 months behind schedule.
- Data Release 2 is scheduled for release on March 15 (2.5 months behind schedule).
- Data Release 3 is scheduled for Oct 1, 2004.
 - DR3 must be loaded with final version of data and made available to the collaboration by July 1.
 - It is important to meet this release date.
- Critical elements in making this happen have already been discussed:
 - Limited data model changes
 - Transfer of knowledge
 - Better technical support
 - Approved hardware plan

Available Resources

- SDSS (ARC-funded Expenses)
 - Personnel Support
 - SDSS provides for 2.1 FTEs of support at JHU and 0.5 FTEs of support at Fermilab for data distribution activities.
 - Hardware Funding
 - Funds are provided for hardware procurements to support data distribution activities at Fermilab and CAS development activities at JHU.
 - SDSS-funded hardware expenses to date = ~\$67K
 - Remaining data distribution hardware budget = \$58K
 - Total anticipated hardware cost = \$125K
 - Does not include value of hardware provided at no charge by HP / Compaq and other sources.

Available Resources (2)

- CAS / SkyServer Development

Also includes effort to host public copies of EDR and DR1

- Johns Hopkins University

- 2.1 FTEs of ARC-funded support

- CAS development
 - Modest SkyServer development
 - Systems administration support

- ~1.5 FTEs of in-kind support for CAS / SkyServer development

- Microsoft Bay Area Research Center

- ~0.1 FTE of in-kind support provided by J. Gray for CAS and SkyServer development

- Princeton University

- ~0.1 FTE of ARC-funded support for R. Lupton to support CAS testing efforts

- Total Development Effort

- 2.1 FTEs of ARC-funded effort
 - 1.7 FTEs of in-kind effort
 - 3.8 FTEs total

Available Resources (3)

- CAS / SkyServer Production Operations
Database loading and hosting at Fermilab
 - Fermilab
 - 0.5 FTEs of ARC-funded support
 - CAS loading and systems admin support
 - ~1.2 FTEs of in-kind support for database loading and hosting
 - 1.1 FTEs of in-kind support from EAG
 - 0.1 FTEs of in-kind support from CD/CSS/CSI
 - New Mexico State University
 - ~0.1 FTE of in-kind support for S. Allam to support CAS testing efforts.
 - Total Operations Effort
 - 0.5 FTEs of ARC-funded effort
 - 1.3 FTEs of in-kind effort
 - 1.8 FTEs total

Resource Availability vs. Requirements

- We have 5.6 FTEs supporting the CAS development, loading, and hosting efforts.
 - 3.8 FTEs on development side
 - 1.8 FTEs on operations side (loading and hosting).
 - All but one FTE work part-time on the CAS effort.
 - Changing data models create stress across the system.
 - Lack of Windows / SQL Server experience on operations side has exacerbated stress levels.
- Ideally we would have at least an additional 0.5 FTE of experienced support on the loading and hosting side (operations).
 - Ideally, the additional 0.5 FTE is added to an existing 0.5 FTE to provide 1.0 FTE of full-time support for loading databases, moving data around, configuring new machines, etc.
- Getting additional help from CD/CSS/CSI to maintain and support the database hardware system will be a huge help.

Summary

- We are resource limited, especially on the operations side. I will be looking into the possibility of adding additional support to this area.
- Controlling data model changes is essential in effectively transferring knowledge of the CAS from development to operations group.
- DR3 provides us with an opportunity to effectively transfer the loading operation to Fermilab, if we manage it well.
- Geographical dispersion has worked against us. We will spend more face time together to facilitate knowledge transfer.
- We would benefit from a written integration plan and a detailed work plan and schedule for future data releases.